Linux Plumbers Conference 2024



Contribution ID: 121

Type: not specified

IO Page Fault for all

The PCI ATS Extended Capability allows peripheral devices to participate in the caching of translations when operating under an IOMMU. Further, the ATS Page Request Interface (PRI) Extension allows devices to handle missing mappings. Currently, PRI is mainly used in the context of Shared Virtual Addressing, requiring support for the Process Address Space Identifier (PASID) capability, but other use cases such as enabling user-space driver driven device verification and reducing memory pinning exists. In this talk we describe how IOMMUFD may be extended in a non-SVA and non-nested context to enable user space processes to handle page requests from VFIO PCI attached devices.

We describe our proposed changes to IOMMUFD and present a user space reference implementation within the libvfn library. In combination with QEMU ATS/PRI emulation, this functionality enables use-case experimentation, hardware/software co-design and device verification for these features. We examine in detail how ATS/PRI is tested, offering insights into the potential of using user space testing frameworks in device validation methodologies.

Primary authors: GRANADOS, Joel; Mr JENSEN, Klaus **Session Classification:** VFIO/IOMMU/PCI MC

Track Classification: VFIO/IOMMU/PCI MC